

Adaptive Low-Complexity Sequential Inference for Dirichlet Process Mixture Models

Theodoros Tsiligkaridis, Keith W. Forsythe
MIT Lincoln Laboratory
Lexington, MA 02141, USA
ttsili@ll.mit.edu
forsythe@ll.mit.edu

September 15, 2015

Abstract

We develop a sequential low-complexity inference procedure for Dirichlet process mixtures of Gaussians for online clustering and parameter estimation when the number of clusters are unknown a-priori. We present an easily computable, closed form parametric expression for the conditional likelihood, in which hyperparameters are recursively updated as a function of the streaming data assuming conjugate priors. Motivated by large-sample asymptotics, we propose a novel adaptive low-complexity design for the Dirichlet process concentration parameter and show that the number of classes grow at most at a logarithmic rate. We further prove that in the large-sample limit, the conditional likelihood and data predictive distribution become asymptotically Gaussian. We demonstrate through experiments on synthetic and real data sets that our approach is superior to other online state-of-the-art methods.

1 Introduction

Dirichlet process mixture models (DPMM) have been widely used for clustering data [9, 11]. Traditional finite mixture models often suffer from overfitting or underfitting of data due to possible mismatch between the model complexity and amount of data. Thus, model selection or model

This work is sponsored by the Assistant Secretary of Defense for Research & Engineering under Air Force Contract #FA8721-05-C-0002. Opinions, interpretations, conclusions and recommendations are those of the author and are not necessarily endorsed by the United States Government.

averaging is required to find the correct number of clusters or the model with the appropriate complexity. This requires significant computation for high-dimensional data sets or large samples. Bayesian nonparametric modeling are alternative approaches to parametric modeling, an example being DPMM's which can automatically infer the number of clusters from the data via Bayesian inference techniques.

The use of Markov chain Monte Carlo (MCMC) methods for Dirichlet process mixtures has made inference tractable [10]. However, these methods can exhibit slow convergence and their convergence can be tough to detect. Alternatives include variational methods [3], which are deterministic algorithms that convert inference to optimization. These approaches can take a significant computational effort even for moderate sized data sets. For large-scale data sets and low-latency applications with streaming data, there is a need for inference algorithms that are much faster and do not require multiple passes through the data. In this work, we focus on low-complexity algorithms that adapt to each sample as they arrive, making them highly scalable. An online algorithm for learning DPMM's based on a sequential variational approximation (SVA) was proposed in [8], and the authors in [15] recently proposed a sequential maximum *a-posterior* (MAP) estimator for the class labels given streaming data. The algorithm is called sequential updating and greedy search (SUGS) and each iteration is composed of a greedy selection step and a posterior update step.

The choice of concentration parameter α is critical for DPMM's as it controls the number of clusters [1]. While most fast DPMM algorithms use a fixed α [6, 4, 7], imposing a prior distribution on α and sampling from it provides more flexibility, but this approach still heavily relies on experimentation and prior knowledge. Thus, many fast inference methods for Dirichlet process mixture models have been proposed that can adapt α to the data, including the works [5] where learning of α is incorporated in the Gibbs sampling analysis, [3] where a Gamma prior is used in a conjugate manner directly in the variational inference algorithm. [15] also account for model uncertainty on the concentration parameter α in a Bayesian manner directly in the sequential inference procedure. This approach can be computationally expensive, as discretization of the domain of α is needed, and its stability highly depends on the initial distribution on α and on the range of values of α . To the best of our knowledge, we are the first to analytically study the evolution and stability of the adapted sequence of α 's in the online learning setting.

In this paper, we propose an adaptive non-Bayesian approach for adapting α motivated by large-sample asymptotics, and call the resulting algorithm ASUGS (Adaptive SUGS). While the basic idea behind ASUGS is directly related to the greedy approach of SUGS, the main contribution is a novel low-complexity stable method for choosing the concentration parameter adaptively as new data arrive, which greatly improves the clustering per-

formance. We derive an upper bound on the number of classes, logarithmic in the number of samples, and further prove that the sequence of concentration parameters that results from this adaptive design is almost bounded. We finally prove, that the conditional likelihood, which is the primary tool used for Bayesian-based online clustering, is asymptotically Gaussian in the large-sample limit, implying that the clustering part of ASUGS asymptotically behaves as a Gaussian classifier. Experiments show that our method outperforms other state-of-the-art methods for online learning of DPMM's.

The paper is organized as follows. In Section 2, we review the sequential inference framework for DPMM's that we will build upon, introduce notation and propose our adaptive modification. In Section 3, the probabilistic data model is given and sequential inference steps are shown. Section 4 contains the growth rate analysis of the number of classes and the adaptively-designed concentration parameters, and Section 5 contains the Gaussian large-sample approximation to the conditional likelihood. Experimental results are shown in Section 6 and we conclude in Section 7.

2 Sequential Inference Framework for DPMM

Here, we review the SUGS framework of [15] for online clustering. Here, the nonparametric nature of the Dirichlet process manifests itself as modeling mixture models with countably infinite components. Let the observations be given by $\mathbf{y}_i \in \mathbb{R}^d$, and γ_i to denote the class label of the i th observation (a latent variable). We define the available information at time i as $\mathbf{y}^{(i)} = \{\mathbf{y}_1, \dots, \mathbf{y}_i\}$ and $\gamma^{(i-1)} = \{\gamma_1, \dots, \gamma_{i-1}\}$. The online sequential updating and greedy search (SUGS) algorithm is summarized next for completeness. Set $\gamma_1 = 1$ and calculate $\pi(\theta_1|\mathbf{y}_1, \gamma_1)$. For $i \geq 2$,

1. Choose best class label for \mathbf{y}_i :

$$\gamma_i \in \arg \max_{1 \leq h \leq k_{i-1}+1} P(\gamma_i = h|\mathbf{y}^{(i)}, \gamma^{(i-1)}).$$

2. Update the posterior distribution using \mathbf{y}_i, γ_i :

$$\pi(\theta_{\gamma_i}|\mathbf{y}^{(i)}, \gamma^{(i)}) \propto f(\mathbf{y}_i|\theta_{\gamma_i})\pi(\theta_{\gamma_i}|\mathbf{y}^{(i-1)}, \gamma^{(i-1)}).$$

where θ_h are the parameters of class h , $f(\mathbf{y}_i|\theta_h)$ is the observation density conditioned on class h and k_{i-1} is the number of classes created at time $i-1$. The algorithm sequentially allocates observations \mathbf{y}_i to classes based on maximizing the conditional posterior probability.

To calculate the posterior probability $P(\gamma_i = h|\mathbf{y}^{(i)}, \gamma^{(i-1)})$, define the variables:

$$L_{i,h}(\mathbf{y}_i) \stackrel{\text{def}}{=} P(\mathbf{y}_i|\gamma_i = h, \mathbf{y}^{(i-1)}, \gamma^{(i-1)}), \quad \pi_{i,h}(\alpha) \stackrel{\text{def}}{=} P(\gamma_i = h|\alpha, \mathbf{y}^{(i-1)}, \gamma^{(i-1)})$$

From Bayes' rule, $P(\gamma_i = h | \mathbf{y}^{(i)}, \gamma^{(i-1)}) \propto L_{i,h}(\mathbf{y}_i) \pi_{i,h}(\alpha)$ for $h = 1, \dots, k_{i-1} + 1$. Here, α is considered fixed at this iteration, and is not updated in a fully Bayesian manner.

According to the Dirichlet process prediction, the predictive probability of assigning observation \mathbf{y}_i to a class h is:

$$\pi_{i,h}(\alpha) = \begin{cases} \frac{m_{i-1}(h)}{i-1+\alpha}, & h = 1, \dots, k_{i-1} \\ \frac{\alpha}{i-1+\alpha}, & h = k_{i-1} + 1 \end{cases} \quad (1)$$

where $m_{i-1}(h) = \sum_{l=1}^{i-1} I(\gamma_l = h)$ counts the number of observations labeled as class h at time $i-1$, and $\alpha > 0$ is the concentration parameter.

2.1 Adaptation of Concentration Parameter α

It is well known that the concentration parameter α has a strong influence on the growth of the number of classes [1]. Our experiments show that in this sequential framework, the choice of α is even more critical. Choosing a fixed α as in the online SVA algorithm of [8] requires cross-validation, which is computationally prohibitive for large-scale data sets. Furthermore, in the streaming data setting where no estimate on the data complexity exists, it is impractical to perform cross-validation. Although the parameter α is handled from a fully Bayesian treatment in [15], a pre-specified grid of possible values α can take, say $\{\alpha_l\}_{l=1}^L$, along with the prior distribution over them, needs to be chosen in advance. Storage and updating of a matrix of size $(k_{i-1} + 1) \times L$ and further marginalization is needed to compute $P(\gamma_i = h | \mathbf{y}^{(i)}, \gamma^{(i-1)})$ at each iteration i . Thus, we propose an alternative data-driven method for choosing α that works well in practice, is simple to compute and has theoretical guarantees.

The idea is to start with a prior distribution on α that favors small α and shape it into a posterior distribution using the data. Define $p_i(\alpha) = p(\alpha | \mathbf{y}^{(i)}, \gamma^{(i)})$ as the posterior distribution formed at time i , which will be used in ASUGS at time $i+1$. Let $p_1(\alpha) \equiv p_1(\alpha | \mathbf{y}^{(1)}, \gamma^{(1)})$ denote the prior for α , e.g., an exponential distribution $p_1(\alpha) = \lambda e^{-\lambda\alpha}$. The dependence on $\mathbf{y}^{(i)}$ and $\gamma^{(i)}$ is trivial only at this first step. Then, by Bayes rule, $p_i(\alpha) \propto p(\mathbf{y}_i, \gamma_i | \mathbf{y}^{(i-1)}, \gamma^{(i-1)}, \alpha) p(\alpha | \mathbf{y}^{(i-1)}, \gamma^{(i-1)}) \propto p_{i-1}(\alpha) \pi_{i,\gamma_i}(\alpha)$ where $\pi_{i,\gamma_i}(\alpha)$ is given in (1). Once this update is made after the selection of γ_i , the α to be used in the next selection step is the mean of the distribution $p_i(\alpha)$, i.e., $\alpha_i = \mathbb{E}[\alpha | \mathbf{y}^{(i)}, \gamma^{(i)}]$. As will be shown in Section 5, the distribution $p_i(\alpha)$ can be approximated by a Gamma distribution with shape parameter k_i and rate parameter $\lambda + \log i$. Under this approximation, we have $\alpha_i = \frac{k_i}{\lambda + \log i}$, only requiring storage and update of one scalar parameter k_i at each iteration i .

The ASUGS algorithm is summarized in Algorithm 1. The selection step may be implemented by sampling the probability mass function $\{q_h^{(i)}\}$. The posterior update step can be efficiently performed by updating the

Algorithm 1 Adaptive Sequential Updating and Greedy Search (ASUGS)

Input: streaming data $\{\mathbf{y}_i\}_{i=1}^{\infty}$, rate parameter $\lambda > 0$.

Set $\gamma_1 = 1$ and $k_1 = 1$. Calculate $\pi(\theta_1|\mathbf{y}_1, \gamma_1)$.

for $i \geq 2$: **do**

(a) Update concentration parameter:

$$\alpha_{i-1} = \frac{k_{i-1}}{\lambda + \log(i-1)}.$$

(b) Choose best label for \mathbf{y}_i :

$$\gamma_i \sim \{q_h^{(i)}\} = \left\{ \frac{L_{i,h}(\mathbf{y}_i)\pi_{i,h}(\alpha_{i-1})}{\sum_{h'} L_{i,h'}(\mathbf{y}_i)\pi_{i,h'}(\alpha_{i-1})} \right\}.$$

(c) Update posterior distribution:

$$\pi(\theta_{\gamma_i}|\mathbf{y}^{(i)}, \gamma^{(i)}) \propto f(\mathbf{y}_i|\theta_{\gamma_i})\pi(\theta_{\gamma_i}|\mathbf{y}^{(i-1)}, \gamma^{(i-1)}).$$

end for

hyperparameters as a function of the streaming data for the case of conjugate distributions. Section 3 derives these updates for the case of multivariate Gaussian observations and conjugate priors for the parameters.

3 Sequential Inference under Unknown Mean & Unknown Covariance

We consider the general case of an unknown mean and covariance for each class. The probabilistic model for the parameters of each class is given as:

$$\mathbf{y}_i|\boldsymbol{\mu}, \mathbf{T} \sim \mathcal{N}(\cdot|\boldsymbol{\mu}, \mathbf{T}), \quad \boldsymbol{\mu}|\mathbf{T} \sim \mathcal{N}(\cdot|\boldsymbol{\mu}_0, c_o\mathbf{T}), \quad \mathbf{T} \sim \mathcal{W}(\cdot|\delta_0, \mathbf{V}_0) \quad (2)$$

where $\mathcal{N}(\cdot|\boldsymbol{\mu}, \mathbf{T})$ denotes the multivariate normal distribution with mean $\boldsymbol{\mu}$ and precision matrix \mathbf{T} , and $\mathcal{W}(\cdot|\delta, \mathbf{V})$ is the Wishart distribution with 2δ degrees of freedom and scale matrix \mathbf{V} . The parameters $\theta = (\boldsymbol{\mu}, \mathbf{T}) \in \mathbb{R}^d \times S_{++}^d$ follow a normal-Wishart joint distribution. The model (16) leads to closed-form expressions for $L_{i,h}(\mathbf{y}_i)$'s due to conjugacy [14].

To calculate the class posteriors, the conditional likelihoods of \mathbf{y}_i given assignment to class h and the previous class assignments need to be calculated first. The conditional likelihood of \mathbf{y}_i given assignment to class h and the history $(\mathbf{y}^{(i-1)}, \gamma^{(i-1)})$ is given by:

$$L_{i,h}(\mathbf{y}_i) = \int f(\mathbf{y}_i|\theta_h)\pi(\theta_h|\mathbf{y}^{(i-1)}, \gamma^{(i-1)})d\theta_h \quad (3)$$

Due to the conjugacy of the distributions, the posterior $\pi(\theta_h|\mathbf{y}^{(i-1)}, \gamma^{(i-1)})$ always has the form:

$$\pi(\theta_h|\mathbf{y}^{(i-1)}, \gamma^{(i-1)}) = \mathcal{N}(\boldsymbol{\mu}_h|\boldsymbol{\mu}_h^{(i-1)}, c_h^{(i-1)}\mathbf{T}_h)\mathcal{W}(\mathbf{T}_h|\delta_h^{(i-1)}, \mathbf{V}_h^{(i-1)})$$

where $\boldsymbol{\mu}_h^{(i-1)}, c_h^{(i-1)}, \delta_h^{(i-1)}, \mathbf{V}_h^{(i-1)}$ are hyperparameters that can be recursively computed as new samples come in. The form of this recursive computation of the hyperparameters is derived in Appendix A. For ease of interpretation and numerical stability, we define $\boldsymbol{\Sigma}_h^{(i)} := \frac{(\mathbf{V}_h^{(i)})^{-1}}{2\delta_h^{(i)}}$ as the inverse of the mean of the Wishart distribution $\mathcal{W}(\cdot|\delta_h^{(i)}, \mathbf{V}_h^{(i)})$. The matrix $\boldsymbol{\Sigma}_h^{(i)}$ has the natural interpretation as the covariance matrix of class h at iteration i . Once the γ_i th component is chosen, the parameter updates for the γ_i th class become:

$$\boldsymbol{\mu}_{\gamma_i}^{(i)} = \frac{1}{1 + c_{\gamma_i}^{(i-1)}}\mathbf{y}_i + \frac{c_{\gamma_i}^{(i-1)}}{1 + c_{\gamma_i}^{(i-1)}}\boldsymbol{\mu}_{\gamma_i}^{(i-1)} \quad (4)$$

$$c_{\gamma_i}^{(i)} = c_{\gamma_i}^{(i-1)} + 1 \quad (5)$$

$$\boldsymbol{\Sigma}_{\gamma_i}^{(i)} = \frac{2\delta_{\gamma_i}^{(i-1)}}{1 + 2\delta_{\gamma_i}^{(i-1)}}\boldsymbol{\Sigma}_{\gamma_i}^{(i-1)} + \frac{1}{1 + 2\delta_{\gamma_i}^{(i-1)}}\frac{c_{\gamma_i}^{(i-1)}}{1 + c_{\gamma_i}^{(i-1)}}(\mathbf{y}_i - \boldsymbol{\mu}_{\gamma_i}^{(i-1)})(\mathbf{y}_i - \boldsymbol{\mu}_{\gamma_i}^{(i-1)})^T \quad (6)$$

$$\delta_{\gamma_i}^{(i)} = \delta_{\gamma_i}^{(i-1)} + \frac{1}{2} \quad (7)$$

If the starting matrix $\boldsymbol{\Sigma}_h^{(0)}$ is positive definite, then all the matrices $\{\boldsymbol{\Sigma}_h^{(i)}\}$ will remain positive definite. Let us return to the calculation of the conditional likelihood (17). By iterated integration, it follows that:

$$L_{i,h}(\mathbf{y}_i) \propto \left(\frac{r_h^{(i-1)}}{2\delta_h^{(i-1)}}\right)^{d/2} \frac{\rho_d(\delta_h^{(i-1)}) \det(\boldsymbol{\Sigma}_h^{(i-1)})^{-1/2}}{\left(1 + \frac{r_h^{(i-1)}}{2\delta_h^{(i-1)}}(\mathbf{y}_i - \boldsymbol{\mu}_h^{(i-1)})^T(\boldsymbol{\Sigma}_h^{(i-1)})^{-1}(\mathbf{y}_i - \boldsymbol{\mu}_h^{(i-1)})\right)^{\delta_h^{(i-1)} + \frac{1}{2}}} \quad (8)$$

where $\rho_d(a) \stackrel{\text{def}}{=} \frac{\Gamma(a+\frac{1}{2})}{\Gamma(a+\frac{1-d}{2})}$ and $r_h^{(i-1)} \stackrel{\text{def}}{=} \frac{c_h^{(i-1)}}{1+c_h^{(i-1)}}$. A detailed mathematical derivation of this conditional likelihood is included in Appendix B. We remark that for the new class $h = k_{i-1} + 1$, $L_{i,k_{i-1}+1}$ has the form (22) with the initial choice of hyperparameters $r^{(0)}, \delta^{(0)}, \boldsymbol{\mu}^{(0)}, \boldsymbol{\Sigma}^{(0)}$.

4 Growth Rate Analysis of Number of Classes & Stability

In this section, we derive a model for the posterior distribution $p_n(\alpha)$ using large-sample approximations, which will allow us to derive growth rates on

the number of classes and the sequence of concentration parameters, showing that the number of classes grows as $\mathbb{E}[k_n] = O(\log^{1+\epsilon} n)$ for ϵ arbitrarily small under certain mild conditions.

The probability density of the α parameter is updated at the j th step in the following fashion:

$$p_{j+1}(\alpha) \propto p_j(\alpha) \cdot \begin{cases} \frac{\alpha}{j+\alpha} & \text{innovation class chosen} \\ \frac{1}{j+\alpha} & \text{otherwise} \end{cases},$$

where only the α -dependent factors in the update are shown. The α -independent factors are absorbed by the normalization to a probability density. Choosing the innovation class pushes mass toward infinity while choosing any other class pushes mass toward zero. Thus there is a possibility that the innovation probability grows in an undesired manner. We assess the growth of the number of innovations $r_n \stackrel{\text{def}}{=} k_n - 1$ under simple assumptions on some likelihood functions that appear naturally in the ASUGS algorithm.

Assuming that the initial distribution of α is $p_1(\alpha) = \lambda e^{-\lambda\alpha}$, the distribution used at step $n + 1$ is proportional to $\alpha^{r_n} \prod_{j=1}^{n-1} (1 + \frac{\alpha}{j})^{-1} e^{-\lambda\alpha}$. We make use of the limiting relation

Theorem 1. *The following asymptotic behavior holds:*

$$\lim_{n \rightarrow \infty} \frac{\log \prod_{j=1}^{n-1} (1 + \frac{\alpha}{j})}{\alpha \log n} = 1.$$

Proof. See Appendix C. □

Using Theorem 1, a large-sample model for $p_n(\alpha)$ is $\alpha^{r_n} e^{-(\lambda + \log n)\alpha}$, suitably normalized. Recognizing this as the Gamma distribution with shape parameter $r_n + 1$ and rate parameter $\lambda + \log n$, its mean is given by $\alpha_n = \frac{r_n + 1}{\lambda + \log n}$. We use the mean in this form to choose class membership in Alg. 1. This asymptotic approximation leads to a very simple scalar update of the concentration parameter; there is no need for discretization for tracking the evolution of continuous probability distributions on α . In our experiments, this approximation is very accurate.

Recall that the innovation class is labeled $K_+ = k_{n-1} + 1$ at the n^{th} step. The modeled updates randomly select a previous class or innovation (new class) by sampling from the probability distribution $\{q_k^{(n)} = P(\gamma_n = k | \mathbf{y}^{(n)}, \gamma^{(n-1)})\}_{k=1}^{K_+}$. Note that $n - 1 = \sum_{k \neq K_+} m_n(k)$, where $m_n(k)$ represents the number of members in class k at time n .

We assume the data follows the Gaussian mixture distribution:

$$p_T(\mathbf{y}) \stackrel{\text{def}}{=} \sum_{h=1}^K \pi_h \mathcal{N}(\mathbf{y} | \boldsymbol{\mu}_h, \boldsymbol{\Sigma}_h) \quad (9)$$

where π_h are the prior probabilities, and $\boldsymbol{\mu}_h, \boldsymbol{\Sigma}_h$ are the parameters of the Gaussian clusters.

Define the mixture-model probability density function, which plays the role of the predictive distribution:

$$\tilde{L}_{n,K_+}(\mathbf{y}) \stackrel{\text{def}}{=} \sum_{k \neq K_+} \frac{m_{n-1}(k)}{n-1} L_{n,k}(\mathbf{y}), \quad (10)$$

so that the probabilities of choosing a previous class or an innovation (using Eq. 1) are proportional to $\sum_{k \neq K_+} \frac{m_{n-1}(k)}{n-1+\alpha_{n-1}} L_{n,k}(\mathbf{y}_n) = \frac{(n-1)}{n-1+\alpha_{n-1}} \tilde{L}_{n,K_+}(\mathbf{y}_n)$ and $\frac{\alpha_{n-1}}{n-1+\alpha_{n-1}} L_{n,K_+}(\mathbf{y}_n)$, respectively. If τ_{n-1} denotes the innovation probability at step n , then we have

$$\left(\rho_{n-1} \frac{\alpha_{n-1} L_{n,K_+}(\mathbf{y}_n)}{n-1+\alpha_{n-1}}, \rho_{n-1} \frac{(n-1) \tilde{L}_{n,K_+}(\mathbf{y}_n)}{n-1+\alpha_{n-1}} \right) = (\tau_{n-1}, 1 - \tau_{n-1}) \quad (11)$$

for some positive proportionality factor ρ_{n-1} .

Define the likelihood ratio (LR) at the beginning of stage n as ¹:

$$l_n(\mathbf{y}) \stackrel{\text{def}}{=} \frac{L_{n,K_+}(\mathbf{y})}{\tilde{L}_{n,K_+}(\mathbf{y})} \quad (12)$$

Conceptually, the mixture (10) represents a modeled distribution fitting the currently observed data. If all “modes” of the data have been observed, it is reasonable to expect that \tilde{L}_{n,K_+} is a good model for future observations. The LR $l_n(\mathbf{y}_n)$ is not large when the future observations are well-modeled by (10). In fact, we expect $\tilde{L}_{n,K_+} \rightarrow p_T$ as $n \rightarrow \infty$, as discussed in Section 5.

Lemma 1. *The following bound holds:*

$$\tau_{n-1} = \frac{l_n(\mathbf{y}_n) \alpha_{n-1}}{n-1 + l_n(\mathbf{y}_n) \alpha_{n-1}} \leq \min \left(\frac{l_n(\mathbf{y}_n) \alpha_{n-1}}{n-1}, 1 \right).$$

Proof. The result follows directly from (11) after a simple calculation. \square

The innovation random variable r_n is described by the random process associated with the probabilities of transition

$$P(r_{n+1} = k | r_n) = \begin{cases} \tau_n, & k = r_n + 1 \\ 1 - \tau_n, & k = r_n \end{cases}. \quad (13)$$

The expectation of r_n is majorized by the expectation of a similar random process, \bar{r}_n , based on the transition probability $\sigma_n \stackrel{\text{def}}{=} \min(\frac{r_n+1}{a_n}, 1)$ instead

¹Here, $L_0(\cdot) \stackrel{\text{def}}{=} L_{n,K_+}(\cdot)$ is independent of n and only depends on the initial choice of hyperparameters as discussed in Sec. 3.

of τ_n as Appendix D shows, where the random sequence $\{a_n\}$ is given by $l_{n+1}(\mathbf{y}_{n+1})^{-1}n(\lambda + \log n)$. The latter can be described as a modification of a Polya urn process with selection probability σ_n . The asymptotic behavior of r_n and related variables is described in the following theorem.

Theorem 2. *Let τ_n be a sequence of real-valued random variables $0 \leq \tau_n \leq 1$ satisfying $\tau_n \leq \frac{r_n+1}{a_n}$ for $n \geq N$, where $a_n = l_{n+1}(\mathbf{y}_{n+1})^{-1}n(\lambda + \log n)$, and where the nonnegative, integer-valued random variables r_n evolve according to (13). Assume the following for $n \geq N$:*

1. $l_n(\mathbf{y}_n) \leq \zeta \quad (a.s.)$
2. $D(p_T \parallel \tilde{L}_{n,K+}) \leq \delta \quad (a.s.)$

where $D(p \parallel q)$ is the Kullback-Leibler divergence between distributions $p(\cdot)$ and $q(\cdot)$. Then, as $n \rightarrow \infty$,

$$r_n = O_P(\log^{1+\zeta}\sqrt{\delta/2} n), \quad \alpha_n = O_P(\log^\zeta\sqrt{\delta/2} n) \quad (14)$$

Proof. See Appendix E. □

Theorem 2 bounds the growth rate of the mean of the number of class innovations and the concentration parameter α_n in terms of the sample size n and parameter ζ . The bounded LR and bounded KL divergence conditions of Thm. 2 manifest themselves in the rate exponents of (14). The experiments section shows that both of the conditions of Thm. 2 hold for all iterations $n \geq N$ for some $N \in \mathbb{N}$. In fact, assuming the correct clustering, the mixture distribution $\tilde{L}_{n,k_{n-1}+1}$ converges to the true mixture distribution p_T , implying that the number of class innovations grows at most as $O(\log^{1+\epsilon} n)$ and the sequence of concentration parameters is $O(\log^\epsilon n)$, where $\epsilon > 0$ can be arbitrarily small.

5 Asymptotic Normality of Conditional Likelihood

In this section, we derive an asymptotic expression for the conditional likelihood (22) in order to gain insight into the steady-state of the algorithm.

We let π_h denote the true prior probability of class h . Using the bounds of the Gamma function in Theorem 1.6 from [2], it follows that $\lim_{a \rightarrow \infty} \frac{\rho_d(a)}{e^{-d/2}(a-1/2)^{d/2}} =$

1. Under normal convergence conditions of the algorithm (with the pruning and merging steps included), all classes $h = 1, \dots, K$ will be correctly identified and populated with approximately $n_{i-1}(h) \approx \pi_h(i-1)$ observations at time $i-1$. Thus, the conditional class prior for each class h converges to π_h as $i \rightarrow \infty$, in virtue of (14), $\pi_{i,h}(\alpha_{i-1}) = \frac{n_{i-1}(h)}{i-1+\alpha_{i-1}} = \frac{\pi_h}{1 + \frac{O_P(\log^\zeta\sqrt{\delta/2}(i-1))}{i-1}} \xrightarrow{i \rightarrow \infty} \pi_h$.

According to (5), we expect $r_h^{(i-1)} \rightarrow 1$ as $i \rightarrow \infty$ since $c_h^{(i-1)} \sim \pi_h(i-1)$.

Also, we expect $2\delta_h^{(i-1)} \sim \pi_h(i-1)$ as $i \rightarrow \infty$ according to (7). Also, from before, $\rho_d(\delta_h^{(i-1)}) \sim e^{-d/2}(\delta_h^{(i-1)} - 1/2)^{d/2} \sim e^{-d/2}(\pi_h \frac{i-1}{2} - \frac{1}{2})^{d/2}$. The parameter updates (4)-(7) imply $\boldsymbol{\mu}_h^{(i)} \rightarrow \boldsymbol{\mu}_h$ and $\boldsymbol{\Sigma}_h^{(i)} \rightarrow \boldsymbol{\Sigma}_h$ as $i \rightarrow \infty$. This follows from the strong law of large numbers, as the updates are recursive implementations of the sample mean and sample covariance matrix. Thus, the large-sample approximation to the conditional likelihood becomes:

$$L_{i,h}(\mathbf{y}_i) \stackrel{i \rightarrow \infty}{\propto} \frac{\lim_{i \rightarrow \infty} \left(1 + \frac{\pi_h^{-1}}{i-1} (\mathbf{y}_i - \boldsymbol{\mu}_h^{(i-1)})^T (\boldsymbol{\Sigma}_h^{(i-1)})^{-1} (\mathbf{y}_i - \boldsymbol{\mu}_h^{(i-1)}) \right)^{-\frac{i-1}{2\pi_h^{-1}}}}{\lim_{i \rightarrow \infty} \det(\boldsymbol{\Sigma}_h^{(i-1)})^{1/2}} \stackrel{i \rightarrow \infty}{\propto} \frac{e^{-\frac{1}{2}(\mathbf{y}_i - \boldsymbol{\mu}_h)^T \boldsymbol{\Sigma}_h^{-1} (\mathbf{y}_i - \boldsymbol{\mu}_h)}}{\sqrt{\det \boldsymbol{\Sigma}_h}} \quad (15)$$

where we used $\lim_{u \rightarrow \infty} (1 + \frac{c}{u})^u = e^c$. The conditional likelihood (15) corresponds to the multivariate Gaussian distribution with mean $\boldsymbol{\mu}_h$ and covariance matrix $\boldsymbol{\Sigma}_h$. A similar asymptotic normality result was recently obtained in [13] for Gaussian observations with a von Mises prior. The asymptotics $\frac{m_{n-1}(h)}{n-1} \rightarrow \pi_h$, $\boldsymbol{\mu}_h^{(n)} \rightarrow \boldsymbol{\mu}_h$, $\boldsymbol{\Sigma}_h^{(n)} \rightarrow \boldsymbol{\Sigma}_h$, $L_{n,h}(\mathbf{y}) \rightarrow \mathcal{N}(\mathbf{y} | \boldsymbol{\mu}_h, \boldsymbol{\Sigma}_h)$ as $n \rightarrow \infty$ imply that the mixture distribution $\tilde{L}_{n,K+}$ in (10) converges to the true Gaussian mixture distribution p_T of (9). Thus, for any small δ , we expect $D(p_T \parallel \tilde{L}_{n,K+}) \leq \delta$ for all $n \geq N$, validating the assumption of Theorem 2.

5.1 Prune & Merge

It is possible that multiple clusters are similar and classes might be created due to outliers, or due to the particular ordering of the streaming data sequence, as also noted in [8]. These effects can be mitigated by adding a pruning and merging step in the ASUGS algorithm.

The pruning step may be implemented as follows. Define $w_h^{(i)} \stackrel{\text{def}}{=} \sum_{j=1}^i q_h^{(j)}$, i.e., the running sum of the posterior weights. The relative weight of each component at the i th iteration may be computed as $\tilde{w}_h^{(i)} = \frac{w_h^{(i)}}{\sum_k w_k^{(i)}}$. If $\tilde{w}_h^{(i)} < \epsilon_r$, then the component is removed.

The merging can be implemented by merging two clusters k_1 and k_2 , once the ℓ_1 distance between the posteriors over time falls below a threshold ϵ_d . This distance is measured as $d_q(k_1, k_2) = \frac{1}{i} \sum_{j=1}^i |q_{k_1}^{(j)} - q_{k_2}^{(j)}|$. This criterion can be implemented in an online fashion by implementing the distance computation recursively. The sufficient statistics are also merged by taking convex combinations $\mu_{k_1}^{(i)} \leftarrow \alpha^{(i)} \mu_{k_1}^{(i)} + (1 - \alpha^{(i)}) \mu_{k_2}^{(i)}$ and $\boldsymbol{\Sigma}_{k_1}^{(i)} \leftarrow \alpha^{(i)} \boldsymbol{\Sigma}_{k_1}^{(i)} + (1 - \alpha^{(i)}) \boldsymbol{\Sigma}_{k_2}^{(i)}$, and by adding $c_{k_1}^{(i)} \leftarrow c_{k_1}^{(i)} + c_{k_2}^{(i)}$ and $\delta_{k_1}^{(i)} \leftarrow \delta_{k_1}^{(i)} + \delta_{k_2}^{(i)}$.

6 Experiments

We apply the ASUGS learning algorithm to a synthetic 16-class example and to a real data set, to verify the stability and accuracy of our method. The experiments show the value of adaptation of the Dirichlet concentration parameter for online clustering and parameter estimation.

Since it is possible that multiple clusters are similar and classes might be created due to outliers, or due to the particular ordering of the streaming data sequence, we add the pruning and merging step in the ASUGS algorithm as done in [8]. We compare ASUGS and ASUGS-PM with SUGS, SUGS-PM, SVA and SVA-PM proposed in [8], since it was shown in [8] that SVA and SVA-PM outperform the block-based methods that perform iterative updates over the entire data set including Collapsed Gibbs Sampling, MCMC with Split-Merge and Truncation-Free Variational Inference.

6.1 Synthetic Data set

We consider learning the parameters of a 16-class Gaussian mixture each with equal variance of $\sigma^2 = 0.025$. The training set was made up of 500 iid samples, and the test set was made up of 1000 iid samples. The clustering results are shown in Fig. 1(a), showing that the ASUGS-based approaches are more stable than SVA-based algorithms. ASUGS-PM performs best and identifies the correct number of clusters, and their parameters. Fig. 1(b) shows the data log-likelihood on the test set (averaged over 100 Monte Carlo trials), the mean and variance of the number of classes at each iteration. The ASUGS-based approaches achieve a higher log-likelihood than SVA-based approaches asymptotically. Fig. 6.1 provides some numerical verification for the assumptions of Theorem 2. As expected, the predictive likelihood $\tilde{L}_{i,K+}$ (10) converges to the true mixture distribution p_T (9), and the likelihood ratio $l_i(\mathbf{y}_i)$ is bounded after enough samples are processed.

6.2 Real Data Set

We applied the online nonparametric Bayesian methods for clustering image data. We used the MNIST data set, which consists of 60,000 training samples, and 10,000 test samples. Each sample is a 28×28 image of a hand-written digit (total of 784 dimensions), and we perform PCA pre-processing to reduce dimensionality to $d = 50$ dimensions as in [7].

We use only a random 1.667% subset, consisting of 1000 random samples for training. This training set contains data from all 10 digits with an approximately uniform proportion. Fig. 3 shows the predictive log-likelihood over the test set, and the mean images for clusters obtained using ASUGS-PM and SVA-PM, respectively. We note that ASUGS-PM achieves higher log-likelihood values and finds all digits correctly using only 23 clusters,




Figure 1: (a) Clustering performance of SVA, SVA-PM, ASUGS and ASUGS-PM on synthetic data set. ASUGS-PM identifies the 16 clusters correctly. (b) Joint log-likelihood on synthetic data, mean and variance of number of classes as a function of iteration. The likelihood values were evaluated on a held-out set of 1000 samples. ASUGS-PM achieves the highest log-likelihood and has the lowest asymptotic variance on the number of classes.

while SVA-PM finds some digits using 56 clusters. Furthermore, the SVA-PM results in noisy-looking image clusters, while ASUGS-PM consistently has clear digits.




Figure 2: Likelihood ratio $l_i(\mathbf{y}_i) = \frac{L_{i,K+}(\mathbf{y}_i)}{\tilde{L}_{i,K+}(\mathbf{y}_i)}$ (left) and L_2 -distance between $\tilde{L}_{i,K+}(\cdot)$ and true mixture distribution p_T (right) for synthetic example (see 1).

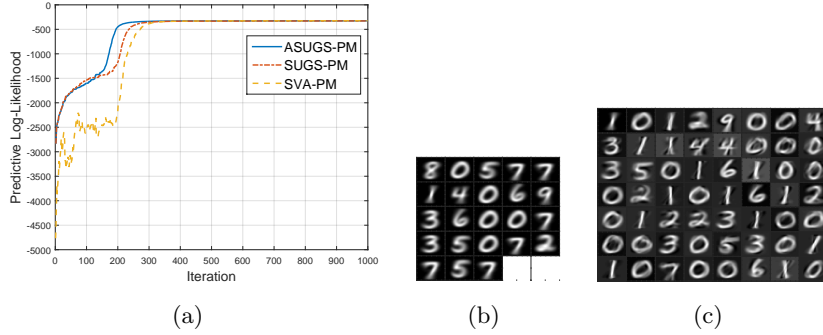


Figure 3: Predictive log-likelihood (a) on test set, mean images for clusters found using ASUGS-PM (b) and SVA-PM (c) on MNIST data set.

6.3 Discussion

Although both SVA and ASUGS methods have similar computational complexity and use decisions and information obtained from processing previous samples in order to decide on class innovations, the mechanics of these methods are quite different. ASUGS uses an adaptive α motivated by asymptotic theory, while SVA uses a fixed α . Furthermore, SVA updates the parameters of all the components at each iteration (in a weighted fashion) while ASUGS only updates the parameters of the most-likely cluster, thus minimizing leakage to unrelated components. The λ parameter of ASUGS does not affect performance as much as the threshold parameter ϵ of SVA does, which often leads to instability requiring lots of pruning and merging steps and increasing latency. This is critical for large data sets or streaming applications, because cross-validation would be required to set ϵ appropriately. We observe higher log-likelihoods and better numerical stability for ASUGS-based

methods in comparison to SVA. The mathematical formulation of ASUGS allows for theoretical guarantees (Theorem 2), and asymptotically normal predictive distribution.

7 Conclusion

We developed a fast online clustering and parameter estimation algorithm for Dirichlet process mixtures of Gaussians, capable of learning in a single data pass. Motivated by large-sample asymptotics, we proposed a novel low-complexity data-driven adaptive design for the concentration parameter and showed it leads to logarithmic growth rates on the number of classes. Through experiments on synthetic and real data sets, we show our method achieves better performance and is as fast as other state-of-the-art online learning DPMM methods.

A Appendix A

We consider the general case of an unknown mean and covariance for each class. Let \mathbf{T} denote the precision (or inverse covariance) matrix. The probabilistic model for the mean and covariance matrix of each class is given as:

$$\begin{aligned} \mathbf{y}_i | \boldsymbol{\mu}, \mathbf{T} &\sim \mathcal{N}(\cdot | \boldsymbol{\mu}, \mathbf{T}) \\ \boldsymbol{\mu} | \mathbf{T} &\sim \mathcal{N}(\cdot | \boldsymbol{\mu}_0, c_0 \mathbf{T}) \\ \mathbf{T} &\sim \mathcal{W}(\cdot | \delta_0, \mathbf{V}_0) \end{aligned} \quad (16)$$

where $\mathcal{N}(\cdot | \boldsymbol{\mu}, \mathbf{T})$ denote the observation density which is assumed to be multivariate normal with mean $\boldsymbol{\mu}$ and precision matrix \mathbf{T} . The parameters $\theta = (\boldsymbol{\mu}, \mathbf{T}) \in \Omega_1 \times \Omega_2$ follow a normal-Wishart joint distribution. The domains here are $\Omega_1 = \mathbb{R}^d$ and $\Omega_2 = S_{++}^d$ is the positive definite cone. This leads to closed-form expressions for $L_{i,h}(\mathbf{y}_i)$'s due to conjugacy [14]. For concreteness, let us write the distributions of the model (16):

$$\begin{aligned} f(\mathbf{y}_i | \theta) &= p(\mathbf{y}_i | \boldsymbol{\mu}, \mathbf{T}) = \frac{\det(\mathbf{T})^{1/2}}{(2\pi)^{d/2}} \exp\left(-\frac{1}{2}(\mathbf{y}_i - \boldsymbol{\mu})^T \mathbf{T}(\mathbf{y}_i - \boldsymbol{\mu})\right) \\ p(\boldsymbol{\mu} | \mathbf{T}) &= p(\theta_1 | \boldsymbol{\Theta}_2) = \frac{\det(c_0 \mathbf{T})^{1/2}}{(2\pi)^{d/2}} \exp\left(-\frac{c_0}{2}(\boldsymbol{\mu} - \boldsymbol{\mu}_0)^T \mathbf{T}(\boldsymbol{\mu} - \boldsymbol{\mu}_0)\right) \\ p(\mathbf{T}) &= p(\boldsymbol{\Theta}_2) = \frac{\det(\mathbf{V}_0)^{-\delta_0}}{2^{d\delta_0} \Gamma_d(\delta_0)} \det(\mathbf{T})^{\delta_0 - \frac{d+1}{2}} \exp\left(-\frac{1}{2}\text{tr}(\mathbf{V}_0^{-1} \mathbf{T})\right) \end{aligned}$$

where $\Gamma_d(\cdot)$ is the multivariate Gamma function.

To calculate the class posteriors, the conditional likelihoods of \mathbf{y}_i given assignment to class h and the previous class assignments need to be calculated first. We derive closed-form expressions for these quantities in this section under the probabilistic model (16).

The conditional likelihood of \mathbf{y}_i given assignment to class h and the history $(\mathbf{y}^{(i-1)}, \gamma^{(i-1)})$ is given by:

$$L_{i,h}(\mathbf{y}_i) = \int f(\mathbf{y}_i | \theta_h) \pi(\theta_h | \mathbf{y}^{(i-1)}, \gamma^{(i-1)}) d\theta_h \quad (17)$$

We thus need to obtain an expression for the posterior distribution $\pi(\theta_h | \mathbf{y}^{(i-1)}, \gamma^{(i-1)})$. Due to the conjugacy of the distributions involved in (16), the posterior distribution $\pi(\theta_h | \mathbf{y}^{(i-1)}, \gamma^{(i-1)})$ always has the form:

$$\pi(\theta_h | \mathbf{y}^{(i-1)}, \gamma^{(i-1)}) = \mathcal{N}(\boldsymbol{\mu}_h | \boldsymbol{\mu}_h^{(i-1)}, c_h^{(i-1)} \mathbf{T}_h) \mathcal{W}(\mathbf{T}_h | \delta_h^{(i-1)}, \mathbf{V}_h^{(i-1)}) \quad (18)$$

where $\boldsymbol{\mu}_h^{(i-1)}, c_h^{(i-1)}, \delta_h^{(i-1)}, \mathbf{V}_h^{(i-1)}$ are hyperparameters that can be recursively computed as new samples come in. This would greatly simplify the

computational complexity of the second step of the SUGS algorithm. Next, we derive the form of this recursive computation of the hyperparameters.

For simplicity of the derivation, let us consider the initial case $\mathbf{y} = \mathbf{y}_1$. Then, from Bayes' rule:

$$p(\theta|\mathbf{y}) = p(\boldsymbol{\mu}, \mathbf{T}|\mathbf{y}) = p(\boldsymbol{\mu}|\mathbf{T}, \mathbf{y})p(\mathbf{T}|\mathbf{y})$$

A.1 Calculation of $p(\boldsymbol{\mu}|\mathbf{T}, \mathbf{y})$

Note the factorization:

$$p(\boldsymbol{\mu}|\mathbf{T}, \mathbf{y}) \propto p(\mathbf{y}|\boldsymbol{\mu}, \mathbf{T})p(\boldsymbol{\mu}|\mathbf{T})$$

According to (16), we can write:

$$\begin{aligned}\mathbf{y} &= \boldsymbol{\mu} + \boldsymbol{\Sigma}^{1/2}\boldsymbol{\epsilon} \\ \boldsymbol{\mu} &= \boldsymbol{\mu}_0 + \boldsymbol{\Sigma}_0^{1/2}\boldsymbol{\epsilon}'\end{aligned}$$

where $\boldsymbol{\epsilon} \sim N(0, I)$, $\boldsymbol{\epsilon}' \sim N(0, I)$, $\boldsymbol{\epsilon}$ is independent of $\boldsymbol{\epsilon}'$ and $\boldsymbol{\Sigma} = \mathbf{T}^{-1}$, $\boldsymbol{\Sigma}_0 = (c_0\mathbf{T})^{-1}$. From this, it follows that the conditional density $p(\boldsymbol{\mu}|\mathbf{T}, \mathbf{y})$ is also multivariate normal with mean $\mathbb{E}[\boldsymbol{\mu}|\mathbf{T}, \mathbf{y}]$ and covariance $\text{Cov}(\boldsymbol{\mu}|\mathbf{T}, \mathbf{y})$. Note that:

$$\begin{aligned}\mathbb{E}[\mathbf{y}|\mathbf{T}] &= \boldsymbol{\mu}_0 \\ \text{Cov}(\mathbf{y}|\mathbf{T}) &= \mathbb{E}[\text{Cov}(\mathbf{y}|\boldsymbol{\mu}, \mathbf{T})|\mathbf{T}] + \text{Cov}(\mathbb{E}[\mathbf{y}|\boldsymbol{\mu}, \mathbf{T}]|\mathbf{T}) \\ &= \boldsymbol{\Sigma} + \boldsymbol{\Sigma}_0 = (1 + c_0^{-1})\mathbf{T}^{-1} \\ \text{Cov}(\boldsymbol{\mu}, \mathbf{y}|\mathbf{T}) &= \boldsymbol{\Sigma}_0\end{aligned}$$

Using these facts, we obtain:

$$\begin{aligned}\mathbb{E}[\boldsymbol{\mu}|\mathbf{T}, \mathbf{y}] &= \mathbb{E}[\boldsymbol{\mu}|\mathbf{T}] + \text{Cov}(\boldsymbol{\mu}, \mathbf{y}|\mathbf{T})\text{Cov}(\mathbf{y}|\mathbf{T})^{-1}(\mathbf{y} - \mathbb{E}[\mathbf{y}|\mathbf{T}]) \\ &= \boldsymbol{\mu}_0 + c_0^{-1}\mathbf{T}^{-1}((1 + c_0^{-1})\mathbf{T}^{-1})^{-1}(\mathbf{y} - \boldsymbol{\mu}_0) \\ &= \boldsymbol{\mu}_0 + c_0^{-1}(1 + c_0^{-1})^{-1}(\mathbf{y} - \boldsymbol{\mu}_0) \\ &= \frac{1}{1 + c_0}\mathbf{y} + \frac{c_0}{1 + c_0}\boldsymbol{\mu}_0 \\ \text{Cov}(\boldsymbol{\mu}|\mathbf{T}, \mathbf{y}) &= \text{Cov}(\boldsymbol{\mu}|\mathbf{T}) - \text{Cov}(\boldsymbol{\mu}, \mathbf{y}|\mathbf{T})\text{Cov}(\mathbf{y}|\mathbf{T})^{-1}\text{Cov}(\mathbf{y}, \boldsymbol{\mu}|\mathbf{T}) \\ &= \boldsymbol{\Sigma}_0 - \boldsymbol{\Sigma}_0(\boldsymbol{\Sigma} + \boldsymbol{\Sigma}_0)^{-1}\boldsymbol{\Sigma}_0^T \\ &= c_0^{-1}\left(1 - \frac{c_0^{-1}}{1 + c_0^{-1}}\right)\mathbf{T}^{-1} \\ &= \frac{c_0^{-1}}{1 + c_0^{-1}}\mathbf{T}^{-1}\end{aligned}$$

Thus, we have:

$$p(\boldsymbol{\mu}|\mathbf{T}, \mathbf{y}) = \mathcal{N}\left(\boldsymbol{\mu} \left| \frac{1}{1+c_0}\mathbf{y} + \frac{c_0}{1+c_0}\boldsymbol{\mu}_0, (1+c_0)\mathbf{T} \right. \right)$$

where the conditional precision matrix becomes $(1+c_0)\mathbf{T}$. As a result, once the γ_i th component is chosen in the SUGS selection step, the parameter updates for the γ_i th class become:

$$\begin{aligned}\boldsymbol{\mu}_{\gamma_i}^{(i)} &= \frac{1}{1+c_{\gamma_i}^{(i-1)}}\mathbf{y}_i + \frac{c_{\gamma_i}^{(i-1)}}{1+c_{\gamma_i}^{(i-1)}}\boldsymbol{\mu}_{\gamma_i}^{(i-1)} \\ c_{\gamma_i}^{(i)} &= c_{\gamma_i}^{(i-1)} + 1\end{aligned}$$

A.2 Calculation of $p(\mathbf{T}|\mathbf{y})$

Next, we focus on calculating $p(\mathbf{T}|\mathbf{y}) = \int_{\mathbb{R}^d} p(\mathbf{T}, \boldsymbol{\mu}|\mathbf{y}) d\boldsymbol{\mu}$, where

$$\begin{aligned}p(\mathbf{T}, \boldsymbol{\mu}|\mathbf{y}) &\propto p(\mathbf{y}|\mathbf{T}, \boldsymbol{\mu})p(\boldsymbol{\mu}|\mathbf{T})p(\mathbf{T}) \\ &\propto \det(\mathbf{T})^{(\delta_0+1/2)-\frac{d+1}{2}} \det(\mathbf{T})^{1/2} \exp\left(-\frac{1}{2}\text{tr}(\mathbf{V}_0^{-1}\mathbf{T})\right) \\ &\quad \times \exp\left(-\frac{1}{2}\left[c_0(\boldsymbol{\mu}-\boldsymbol{\mu}_0)^T\mathbf{T}(\boldsymbol{\mu}-\boldsymbol{\mu}_0) + (\mathbf{y}-\boldsymbol{\mu})^T\mathbf{T}(\mathbf{y}-\boldsymbol{\mu})\right]\right)\end{aligned}$$

Rewriting the term inside the brackets by completing the square, we obtain:

$$\begin{aligned}&c_0(\boldsymbol{\mu}-\boldsymbol{\mu}_0)^T\mathbf{T}(\boldsymbol{\mu}-\boldsymbol{\mu}_0) + (\mathbf{y}-\boldsymbol{\mu})^T\mathbf{T}(\mathbf{y}-\boldsymbol{\mu}) \\ &= c_0\|\mathbf{T}^{1/2}\boldsymbol{\mu} - \mathbf{T}^{1/2}\boldsymbol{\mu}_0\|_2^2 + \|\mathbf{T}^{1/2}\mathbf{y} - \mathbf{T}^{1/2}\boldsymbol{\mu}\|_2^2 \\ &= (1+c_0)\left\{\|\mathbf{T}^{1/2}\boldsymbol{\mu}\|_2^2 - 2\left\langle\mathbf{T}^{1/2}\boldsymbol{\mu}, \frac{c_0\mathbf{T}^{1/2}\boldsymbol{\mu}_0 + \mathbf{T}^{1/2}\mathbf{y}}{1+c_0}\right\rangle + \frac{c_0\|\mathbf{T}^{1/2}\boldsymbol{\mu}_0\|_2^2 + \|\mathbf{T}^{1/2}\mathbf{y}\|_2^2}{1+c_0}\right\} \\ &= (1+c_0)\left\{\left\|\mathbf{T}^{1/2}\boldsymbol{\mu} - \frac{c_0\mathbf{T}^{1/2}\boldsymbol{\mu}_0 + \mathbf{T}^{1/2}\mathbf{y}}{1+c_0}\right\|_2^2 - \left\|\frac{c_0\mathbf{T}^{1/2}\boldsymbol{\mu}_0 + \mathbf{T}^{1/2}\mathbf{y}}{1+c_0}\right\|_2^2 + \frac{c_0\|\mathbf{T}^{1/2}\boldsymbol{\mu}_0\|_2^2 + \|\mathbf{T}^{1/2}\mathbf{y}\|_2^2}{1+c_0}\right\}\end{aligned}$$

Integrating out $\boldsymbol{\mu}$, we obtain:

$$\begin{aligned}&\int \exp\left(-\frac{1}{2}\left[c_0(\boldsymbol{\mu}-\boldsymbol{\mu}_0)^T\mathbf{T}(\boldsymbol{\mu}-\boldsymbol{\mu}_0) + (\mathbf{y}-\boldsymbol{\mu})^T\mathbf{T}(\mathbf{y}-\boldsymbol{\mu})\right]\right) d\boldsymbol{\mu} \\ &= \exp\left(-\frac{1+c_0}{2}\left(\frac{c_0\|\mathbf{T}^{1/2}\boldsymbol{\mu}_0\|_2^2 + \|\mathbf{T}^{1/2}\mathbf{y}\|_2^2}{1+c_0} - \left\|\frac{c_0\mathbf{T}^{1/2}\boldsymbol{\mu}_0 + \mathbf{T}^{1/2}\mathbf{y}}{1+c_0}\right\|_2^2\right)\right) \\ &\quad \times \int \exp\left(-\frac{1}{2}\left\|\mathbf{T}^{1/2}\boldsymbol{\mu} - \frac{c_0\mathbf{T}^{1/2}\boldsymbol{\mu}_0 + \mathbf{T}^{1/2}\mathbf{y}}{1+c_0}\right\|_2^2\right) d\boldsymbol{\mu} \\ &\propto \det(\mathbf{T})^{-1/2} \exp\left(-\frac{1}{2}\frac{c_0}{1+c_0}(\mathbf{y}-\boldsymbol{\mu}_0)^T\mathbf{T}(\mathbf{y}-\boldsymbol{\mu}_0)\right)\end{aligned}$$

Using this result, we obtain:

$$p(\mathbf{T}|\mathbf{y}) \propto \det(\mathbf{T})^{(\delta_0+1/2)-\frac{d+1}{2}} \exp\left(-\frac{1}{2}\text{tr}\left(\mathbf{T}\left\{\mathbf{V}_0^{-1} + \frac{c_0}{1+c_0}(\mathbf{y}-\boldsymbol{\mu}_0)(\mathbf{y}-\boldsymbol{\mu}_0)^T\right\}\right)\right)$$

As a result, the conditional density is recognized to be a Wishart distribution

$$\mathcal{W}\left(\mathbf{T} \middle| \delta_0 + \frac{1}{2}, \left\{\mathbf{V}_0^{-1} + \frac{c_0}{1+c_0}(\mathbf{y}-\boldsymbol{\mu}_0)(\mathbf{y}-\boldsymbol{\mu}_0)^T\right\}^{-1}\right).$$

Thus, the parameter updates for the γ_i th class become:

$$\begin{aligned} \delta_{\gamma_i}^{(i)} &= \delta_{\gamma_i}^{(i-1)} + \frac{1}{2} \\ \mathbf{V}_{\gamma_i}^{(i)} &= \left\{(\mathbf{V}_{\gamma_i}^{(i-1)})^{-1} + \frac{c_{\gamma_i}^{(i-1)}}{1+c_{\gamma_i}^{(i-1)}}(\mathbf{y}_i - \boldsymbol{\mu}_{\gamma_i}^{(i-1)})(\mathbf{y}_i - \boldsymbol{\mu}_{\gamma_i}^{(i-1)})^T\right\}^{-1} \end{aligned} \quad (19)$$

For numerical stability and ease of interpretation, we define

$$\boldsymbol{\Sigma}_h^{(i)} := \frac{(\mathbf{V}_h^{(i)})^{-1}}{2\delta_h^{(i)}}.$$

This is the inverse of the mean of the Wishart distribution $\mathcal{W}(\cdot|\delta_h^{(i)}, \mathbf{V}_h^{(i)})$, and can be interpreted as the covariance matrix of class h at iteration i . From (19), we have:

$$\begin{aligned} \boldsymbol{\Sigma}_h^{(i)} &= \frac{(\mathbf{V}_h^{(i)})^{-1}}{2\delta_h^{(i)}} \\ &= \frac{2\delta_h^{(i-1)}}{2\delta_h^{(i)}} \frac{(\mathbf{V}_h^{(i-1)})^{-1}}{2\delta_h^{(i-1)}} + \frac{1}{2\delta_h^{(i)}} \frac{c_h^{(i-1)}}{1+c_h^{(i-1)}}(\mathbf{y}_i - \boldsymbol{\mu}_h^{(i-1)})(\mathbf{y}_i - \boldsymbol{\mu}_h^{(i-1)})^T \\ &= \frac{2\delta_h^{(i-1)}}{1+2\delta_h^{(i-1)}} \boldsymbol{\Sigma}_h^{(i-1)} + \frac{1}{1+2\delta_h^{(i-1)}} \frac{c_h^{(i-1)}}{1+c_h^{(i-1)}}(\mathbf{y}_i - \boldsymbol{\mu}_h^{(i-1)})(\mathbf{y}_i - \boldsymbol{\mu}_h^{(i-1)})^T \end{aligned}$$

Thus, the recursive updates (19) can be equivalently restated as:

$$\begin{aligned} \delta_{\gamma_i}^{(i)} &= \delta_{\gamma_i}^{(i-1)} + \frac{1}{2} \\ \boldsymbol{\Sigma}_h^{(i)} &= \frac{2\delta_h^{(i-1)}}{1+2\delta_h^{(i-1)}} \boldsymbol{\Sigma}_h^{(i-1)} + \frac{1}{1+2\delta_h^{(i-1)}} \frac{c_h^{(i-1)}}{1+c_h^{(i-1)}}(\mathbf{y}_i - \boldsymbol{\mu}_h^{(i-1)})(\mathbf{y}_i - \boldsymbol{\mu}_h^{(i-1)})^T \end{aligned}$$

If the starting matrix $\boldsymbol{\Sigma}_h^{(0)}$ is positive definite, then all the matrices $\{\boldsymbol{\Sigma}_h^{(i)}\}$ will remain positive definite.

B Appendix B

Now, let us return to the calculation of (17).

$$\begin{aligned} L_{i,h}(\mathbf{y}_i) &= \int_{S_{++}^d} \int_{\mathbb{R}^d} \mathcal{N}(\mathbf{y}_i | \boldsymbol{\mu}, \mathbf{T}) \mathcal{N}(\boldsymbol{\mu} | \boldsymbol{\mu}_h^{(i-1)}, c_h^{(i-1)} \mathbf{T}) \mathcal{W}(\mathbf{T} | \delta_h^{(i-1)}, \mathbf{V}_h^{(i-1)}) d\boldsymbol{\mu} d\mathbf{T} \\ &= \int_{S_{++}^d} \mathcal{W}(\mathbf{T} | \delta_h^{(i-1)}, \mathbf{V}_h^{(i-1)}) \left\{ \int_{\mathbb{R}^d} \mathcal{N}(\mathbf{y}_i | \boldsymbol{\mu}, \mathbf{T}) \mathcal{N}(\boldsymbol{\mu} | \boldsymbol{\mu}_h^{(i-1)}, c_h^{(i-1)} \mathbf{T}) d\boldsymbol{\mu} \right\} d\mathbf{T} \end{aligned}$$

Evaluating the inner integral within the brackets:

$$\begin{aligned} &\int_{\mathbb{R}^d} \mathcal{N}(\mathbf{y}_i | \boldsymbol{\mu}, \mathbf{T}) \mathcal{N}(\boldsymbol{\mu} | \boldsymbol{\mu}_h^{(i-1)}, c_h^{(i-1)} \mathbf{T}) d\boldsymbol{\mu} \\ &\propto \det(\mathbf{T})^{1/2} \det(c_h^{(i-1)} \mathbf{T})^{1/2} \\ &\quad \times \int_{\mathbb{R}^d} \exp \left(-\frac{1}{2} \left[c_h^{(i-1)} (\boldsymbol{\mu} - \boldsymbol{\mu}_h^{(i-1)})^T \mathbf{T} (\boldsymbol{\mu} - \boldsymbol{\mu}_h^{(i-1)}) + (\mathbf{y}_i - \boldsymbol{\mu})^T \mathbf{T} (\mathbf{y}_i - \boldsymbol{\mu}) \right] \right) d\boldsymbol{\mu} \\ &= \det(\mathbf{T})^{1/2} \det(c_h^{(i-1)} \mathbf{T})^{1/2} \\ &\quad \times \exp \left(-\frac{1}{2} \frac{c_h^{(i-1)}}{1 + c_h^{(i-1)}} (\mathbf{y}_i - \boldsymbol{\mu}_h^{(i-1)})^T \mathbf{T} (\mathbf{y}_i - \boldsymbol{\mu}_h^{(i-1)}) \right) \\ &\quad \times \int \exp \left(-\frac{1 + c_h^{(i-1)}}{2} (\boldsymbol{\mu} - \mathbf{b})^T \mathbf{T} (\boldsymbol{\mu} - \mathbf{b}) \right) d\boldsymbol{\mu} \\ &\propto \frac{\det(\mathbf{T})^{1/2} \det(c_h^{(i-1)} \mathbf{T})^{1/2}}{\det((1 + c_h^{(i-1)}) \mathbf{T})^{1/2}} \exp \left(-\frac{1}{2} \text{tr} \left(\mathbf{T} \left\{ \frac{c_h^{(i-1)}}{1 + c_h^{(i-1)}} (\mathbf{y}_i - \boldsymbol{\mu}_h^{(i-1)}) (\mathbf{y}_i - \boldsymbol{\mu}_h^{(i-1)})^T \right\} \right) \right) \\ &= \left(\frac{c_h^{(i-1)}}{1 + c_h^{(i-1)}} \right)^{d/2} \det(\mathbf{T})^{1/2} \exp \left(-\frac{1}{2} \text{tr} \left(\mathbf{T} \left\{ \frac{c_h^{(i-1)}}{1 + c_h^{(i-1)}} (\mathbf{y}_i - \boldsymbol{\mu}_h^{(i-1)}) (\mathbf{y}_i - \boldsymbol{\mu}_h^{(i-1)})^T \right\} \right) \right) \end{aligned}$$

Using this closed-form expression for the inner integral, we further obtain:

$$L_{i,h}(\mathbf{y}_i) \propto \left(\frac{c_h^{(i-1)}}{1 + c_h^{(i-1)}} \right)^{d/2} \int_{S_{++}^d} \frac{\det(\mathbf{V}_h^{(i-1)})^{-\delta_h^{(i-1)}}}{2^{d\delta_h^{(i-1)}} \Gamma_d(\delta_h^{(i-1)})} \det(\mathbf{T})^{(\delta_h^{(i-1)}+1/2)-\frac{d+1}{2}} \\ \times \exp \left(-\frac{1}{2} \text{tr} \left(\mathbf{T} \left\{ (\mathbf{V}_h^{(i-1)})^{-1} + \frac{c_h^{(i-1)}}{1 + c_h^{(i-1)}} (\mathbf{y}_i - \boldsymbol{\mu}_h^{(i-1)})(\mathbf{y}_i - \boldsymbol{\mu}_h^{(i-1)})^T \right\} \right) \right) d\mathbf{T} \quad (20)$$

$$\propto \left(\frac{c_h^{(i-1)}}{1 + c_h^{(i-1)}} \right)^{d/2} \frac{\Gamma_d(\delta_h^{(i-1)} + \frac{1}{2})}{\Gamma_d(\delta_h^{(i-1)})} \\ \times \frac{\det(\mathbf{V}_h^{(i-1)})^{-\delta_h^{(i-1)}}}{\det \left(\left\{ (\mathbf{V}_h^{(i-1)})^{-1} + \frac{c_h^{(i-1)}}{1 + c_h^{(i-1)}} (\mathbf{y}_i - \boldsymbol{\mu}_h^{(i-1)})(\mathbf{y}_i - \boldsymbol{\mu}_h^{(i-1)})^T \right\}^{-1} \right)^{-(\delta_h^{(i-1)} + \frac{1}{2})}} \\ = \left(r_h^{(i-1)} \right)^{d/2} \frac{\Gamma_d(\delta_h^{(i-1)} + \frac{1}{2})}{\Gamma_d(\delta_h^{(i-1)})} \frac{\det((\mathbf{V}_h^{(i-1)})^{-1})^{-1/2}}{\det \left(\mathbf{I}_d + r_h^{(i-1)} (\mathbf{y}_i - \boldsymbol{\mu}_h^{(i-1)})(\mathbf{y}_i - \boldsymbol{\mu}_h^{(i-1)})^T \mathbf{V}_h^{(i-1)} \right)^{\delta_h^{(i-1)} + \frac{1}{2}}} \\ = \left(r_h^{(i-1)} \right)^{d/2} \frac{\Gamma_d(\delta_h^{(i-1)} + \frac{1}{2})}{\Gamma_d(\delta_h^{(i-1)})} \frac{\det(\mathbf{V}_h^{(i-1)})^{1/2}}{\left(1 + r_h^{(i-1)} (\mathbf{y}_i - \boldsymbol{\mu}_h^{(i-1)})^T \mathbf{V}_h^{(i-1)} (\mathbf{y}_i - \boldsymbol{\mu}_h^{(i-1)}) \right)^{\delta_h^{(i-1)} + \frac{1}{2}}} \quad (21)$$

$$= \left(\frac{r_h^{(i-1)}}{2\delta_h^{(i-1)}} \right)^{d/2} \frac{\Gamma_d(\delta_h^{(i-1)} + \frac{1}{2})}{\Gamma_d(\delta_h^{(i-1)})} \frac{\det((\boldsymbol{\Sigma}_h^{(i-1)})^{-1})^{1/2}}{\left(1 + \frac{r_h^{(i-1)}}{2\delta_h^{(i-1)}} (\mathbf{y}_i - \boldsymbol{\mu}_h^{(i-1)})^T (\boldsymbol{\Sigma}_h^{(i-1)})^{-1} (\mathbf{y}_i - \boldsymbol{\mu}_h^{(i-1)}) \right)^{\delta_h^{(i-1)} + \frac{1}{2}}} \quad (22)$$

where we used the determinant identity $\det(\mathbf{I} + \mathbf{a}\mathbf{b}^T\mathbf{M}) = 1 + \mathbf{b}^T\mathbf{M}\mathbf{a}$ in the last step. We also defined $r_h^{(i)} := \frac{c_h^{(i)}}{1+c_h^{(i)}}$ and used $\mathbf{V}_h^{(i)} = \frac{(\boldsymbol{\Sigma}_h^{(i)})^{-1}}{2\delta_h^{(i)}}$.

C Appendix C

Proof. It is sufficient to establish the limit for $\lim_{N \rightarrow \infty} \sum_{k=m}^N \log(1+\alpha/k) / \log N$ for fixed m . Choose m such that $|\alpha| < m-1$ and use $\log(1-x) = \sum_{k=1}^{\infty} x^k/k$ for $|x| < 1$ to get

$$\sum_{k=m}^N \log \left(1 + \frac{\alpha}{k} \right) = \sum_{l=1}^{\infty} (-1)^{l+1} \frac{\alpha^l}{l} \sum_{k=m}^N \frac{1}{k^l}. \quad (23)$$

Separate (23) into two terms:

$$\sum_{l=1}^{\infty} (-1)^{l+1} \frac{\alpha^l}{l} \sum_{k=m}^N \frac{1}{k^l} = \alpha \sum_{k=m}^N \frac{1}{k} + \sum_{l=2}^{\infty} (-1)^{l+1} \frac{\alpha^l}{l} \sum_{k=m}^N \frac{1}{k^l}. \quad (24)$$

The first term is expressed in terms of the Euler-Mascheroni constant γ_e as

$$\sum_{k=m}^N \frac{1}{k} = \log N - \gamma_e - \sum_{k=1}^{m-1} \frac{1}{k} + o(1).$$

Thus, dividing by $\log N$ and taking the limit $N \rightarrow \infty$ we have a limiting value of unity. The second term of (24) is bounded. To see this, use, for $l > 1$,

$$\sum_{k=m}^{\infty} \frac{1}{k^l} \leq \int_{m-1}^{\infty} \frac{dx}{x^l} = \frac{1}{l-1} (m-1)^{-(l-1)}.$$

Then the second term of (24) is bounded by

$$\begin{aligned} \sum_{l=2}^{\infty} \frac{\alpha^l}{l} \sum_{k=m}^{\infty} \frac{1}{k^l} &\leq \sum_{l=2}^{\infty} \frac{\alpha^l}{l(l-1)} (m-1)^{-(l-1)} \\ &= (m-1) \sum_{l=2}^{\infty} \frac{1}{l(l-1)} \left(\frac{\alpha}{m-1} \right)^l < \infty. \end{aligned}$$

The result follows since the second term, being bounded, vanished when dividing by $\log N$ and taking the limit $N \rightarrow \infty$. \square

D Appendix D

Lemma 2. *Let r_n and \bar{r}_n be random sequences with the update laws*

$$\begin{aligned} P(r_{n+1} = r_n + 1) &= \tau_n \\ P(r_{n+1} = r_n) &= 1 - \tau_n \end{aligned}$$

and

$$\begin{aligned} P(\bar{r}_{n+1} = \bar{r}_n + 1) &= \sigma_n \\ P(\bar{r}_{n+1} = \bar{r}_n) &= 1 - \sigma_n, \end{aligned}$$

and assume $\sigma_n \geq \tau_n$ for all $n \geq 1$ and that $\bar{r}_0 = r_0 = 0$. Then $\mathbb{E}[\bar{r}_n] \geq \mathbb{E}[r_n]$ for all $n \geq 1$.

Proof. We first use induction to show that $P(r_n > t) \leq P(\bar{r}_n > t)$ holds for all n .

The base case is trivial because $r_0 = \bar{r}_0$. We next prove that given

$$P(r_n > t) \leq P(\bar{r}_n > t) \quad (25)$$

for a particular n and all $t \in \mathbb{N}$, the same inequality holds for $n + 1$. We have

$$\begin{aligned} P(r_{n+1} > t) &= (1 - \tau_n)P(r_n > t) + \tau_n P(r_n > t - 1) \\ &\leq (1 - \tau_n)P(\bar{r}_n > t) + \tau_n P(\bar{r}_n > t - 1) \\ &\leq (1 - \sigma_n)P(\bar{r}_n > t) + \sigma_n P(\bar{r}_n > t - 1) \\ &= P(\bar{r}_{n+1} > t), \end{aligned} \quad (26)$$

where we used the inductive hypothesis (25) and the inequality $P(\bar{r}_n > t) \leq P(\bar{r}_n > t - 1)$. Thus, by induction, the inequality (25) holds for all n . Using (25), we further obtain:

$$\mathbb{E}[r_n] = \int_0^\infty P(r_n > t) dt \leq \int_0^\infty P(\bar{r}_n > t) dt = \mathbb{E}[\bar{r}_n]$$

The proof is complete. \square

E Appendix E

Proof. We can study the generalized Polya urn model in the slightly modified form:

$$P(\bar{r}_{n+1} = k | \bar{r}_n) = \begin{cases} \frac{\bar{r}_n + 1}{a_n}, & \text{if } k = \bar{r}_n + 1 \\ 1 - \frac{\bar{r}_n + 1}{a_n}, & \text{if } k = \bar{r}_n \end{cases} \quad (27)$$

Taking the conditional expectation of \bar{r}_{n+1} with respect to the filtration $\mathcal{F}_{n+1} \stackrel{\text{def}}{=} \sigma(\bar{r}_1, \dots, \bar{r}_n, \gamma_1, \dots, \gamma_{n+1}, \mathbf{y}_1, \dots, \mathbf{y}_{n+1})$, we get $\mathbb{E}[\bar{r}_{n+1} | \mathcal{F}_{n+1}] = (\bar{r}_n + 1) \left(1 + \frac{1}{a_n}\right) - 1$. Set $x_n := \bar{r}_n + 1$. Rewriting this and using the definition of a_n , we obtain:

$$\mathbb{E}[x_{n+1} | \mathcal{F}_{n+1}] \leq x_n \left(1 + \frac{l_{n+1}(\mathbf{y}_{n+1})}{n \log n}\right) \quad (28)$$

Next, we seek an upper bound on the conditional expectation $\mathbb{E}[l_k(\mathbf{y}_k) | \mathcal{F}_{k-1}]$. This quantity can be bounded using convex duality [12]:

$$\mathbb{E}[l_k(\mathbf{y}_k) | \mathcal{F}_{k-1}] \leq 1 + \frac{1}{s} D(p_T \parallel \tilde{L}_{k,K+}) + \frac{1}{s} \log \mathbb{E}_{\tilde{L}_{k,K+}}[e^{s(l_k(\mathbf{y}_k) - 1)}]$$

For $k \geq N$, $l_k(\mathbf{y}_k) \leq \zeta$ and $\mathbb{E}_{\tilde{L}_{k,K+}}[l_k(\mathbf{y}_k)] = 1$. By Hoeffding's inequality, $\mathbb{E}_{\tilde{L}_{k,K+}}[e^{s(l_k(\mathbf{y}_k) - 1)}] \leq e^{s^2 \zeta^2 / 8}$. Using this bound, we obtain for $k \geq N$,

$\mathbb{E}[l_k(\mathbf{y}_k)|\mathcal{F}_{k-1}] \leq 1 + \delta/s + s\zeta^2/8$. Minimizing this as a function of $s > 0$, we obtain:

$$\mathbb{E}[l_k(\mathbf{y}_k)|\mathcal{F}_{k-1}] \leq 1 + \zeta\sqrt{\frac{\delta}{2}} \quad (29)$$

Next, we upper bound $\mathbb{E}[x_{n+1}|\mathcal{F}_N]$ recursively. Taking the conditional expectation of both sides of (28), we obtain:

$$\mathbb{E}[x_{n+1}|\mathcal{F}_n] \leq \mathbb{E}\left[x_n \left(1 + \frac{l_{n+1}(\mathbf{y}_{n+1})}{n \log n}\right) \middle| \mathcal{F}_n\right] \quad (30)$$

We note that the function $l_{n+1}(\cdot)$ is \mathcal{F}_n -measurable. This follows since by definition, $l_{n+1}(\cdot) = \frac{L_0(\cdot)}{\sum_{h=1}^{k_n} \frac{m_n(h)}{n} L_{n+1,h}(\cdot)}$, and $m_n(h) = \sum_{l=1}^n I(\gamma_l = h)$ and $L_{n+1,h}(\cdot)$ are both \mathcal{F}_n -measurable (due to the parameter updates and (22)). Also note that $x_n = \bar{r}_n + 1$ is randomly determined by a biased coin flip given \mathcal{F}_n , increasing by 1 with probability $\frac{x_{n-1}}{a_{n-1}}$ and staying the same with probability $1 - \frac{x_{n-1}}{a_{n-1}}$. Since a_{n-1} is \mathcal{F}_n -measurable, it follows that x_n and $l_{n+1}(\mathbf{y}_{n+1})$ are conditionally independent given the history \mathcal{F}_n . Using this conditional independence, we obtain from (30):

$$\mathbb{E}[x_{n+1}|\mathcal{F}_n] \leq \mathbb{E}[x_n|\mathcal{F}_n] \left(1 + \frac{\mathbb{E}[l_{n+1}(\mathbf{y}_{n+1})|\mathcal{F}_n]}{n \log n}\right) \leq \mathbb{E}[x_n|\mathcal{F}_n] \left(1 + \frac{1 + \zeta\sqrt{\delta/2}}{n \log n}\right) \quad (31)$$

where we used the bound (29) in the last inequality. Repeatedly conditioning and using (28) and (31): $\mathbb{E}[x_{n+1}|\mathcal{F}_N] \leq \prod_{k=N}^n \left(1 + \frac{1 + \zeta\sqrt{\delta/2}}{k \log k}\right) \mathbb{E}[x_N|\mathcal{F}_N] \leq$

$C_0 N \log^{1+\zeta\sqrt{\delta/2}} n$, where we used the Lemma in Appendix F and $C_0 = C(1 + \zeta\sqrt{\delta/2}, N)$, $x_N \leq N$ in the last inequality. Taking the unconditional expectation and using $\mathbb{E}[r_n + 1] \leq \mathbb{E}[\bar{r}_n + 1]$ (see Appendix D) yields the bound $\mathbb{E}[r_n + 1] \leq C_0 N \log^{1+\zeta\sqrt{\delta/2}} n$. Markov's inequality then yields $\mathbb{P}\left(\frac{r_n+1}{C_0 N \log^{1+\zeta\sqrt{\delta/2}} n} > K\right) \leq \frac{1}{K}$ which implies (14) by taking $K \rightarrow \infty$. Since $\alpha_n = \frac{r_n+1}{\lambda + \log n}$, the bound in (14) follows from a similar argument. The proof is complete. \square

F Appendix F

Lemma 3. *The following upper bound holds with constant $C(\phi, N) = e^{\frac{\phi}{N \log N}} / \log^\phi N$:*

$$\prod_{k=N}^n \left(1 + \frac{\phi}{k \log k}\right) \leq C(\phi, N) \log^\phi n$$

Proof. Using the elementary inequality $\log(1+x) \leq x$ for $x > -1$, we obtain:

$$\begin{aligned}
& \log \left(\prod_{k=N}^n \left(1 + \frac{\phi}{k \log k} \right) \right) = \sum_{k=N}^n \log \left(1 + \frac{\phi}{k \log k} \right) \\
& \leq \sum_{k=N}^n \frac{\phi}{k \log k} \leq \phi \left(\int_N^n \frac{dx}{x \log x} + \frac{1}{N \log N} \right) \\
& = \phi \left(\int_{\log N}^{\log n} \frac{dt}{t} + \frac{1}{N \log N} \right) \\
& = \log \left(\frac{\log^\phi n}{\log^\phi N} \right) + \frac{\phi}{N \log N}
\end{aligned}$$

Taking the exponential of both sides yields the desired inequality. \square

References

- [1] C. E. Antoniak, *Mixtures of Dirichlet Processes with Applications to Bayesian Nonparametric Problems*, The Annals of Statistics **2** (1974), no. 6, 1152–1174.
- [2] N. Batir, *Inequalities for the Gamma Function*, Archiv der Mathematik **91** (2008), no. 6, 554–563.
- [3] D. M. Blei and M. I. Jordan, *Variational Inference for Dirichlet Process Mixtures*, Bayesian Analysis **1** (2006), no. 1, 121–144.
- [4] H. Daume, *Fast Search for Dirichlet Process Mixture Models*, Conference on Artificial Intelligence and Statistics, 2007.
- [5] M. D. Escobar and M. West, *Bayesian Density Estimation and Inference using Mixtures*, Journal of the American Statistical Association **90** (1995), no. 430, 577–588.
- [6] P. Fearnhead, *Particle Filters for Mixture Models with an Unknown Number of Components*, Statistics and Computing **14** (2004), 11–21.
- [7] K. Kurihara, M. Welling, and N. Vlassis, *Accelerated Variational Dirichlet Mixture Models*, Advances in Neural Information Processing Systems (NIPS), 2006.
- [8] Dahua Lin, *Online learning of nonparametric mixture models via sequential variational approximation*, Advances in Neural Information Processing Systems 26 (C.J.C. Burges, L. Bottou, M. Welling, Z. Ghahramani, and K.Q. Weinberger, eds.), Curran Associates, Inc., 2013, pp. 395–403.

- [9] R. M. Neal, *Bayesian Mixture Modeling*, Proceedings of the Workshop on Maximum Entropy and Bayesian Methods of Statistical Analysis, vol. 11, 1992, pp. 197–211.
- [10] ———, *Markov chain sampling methods for Dirichlet process mixture models*, Journal of Computational and Graphical Statistics **9** (2000), no. 2, 249–265.
- [11] C. E. Rasmussen, *The infinite gaussian mixture model*, Advances in Neural Information Processing Systems 12, MIT Press, 2000, pp. 554–560.
- [12] Matthias W. Seeger, *Bayesian Gaussian Process Models: PAC-Bayesian Generalization Error Bounds and Sparse Approximations*, Ph.D. thesis, University of Edinburgh, 2003.
- [13] T. Tsiligkaridis and K. W. Forsythe, *A Sequential Bayesian Inference Framework for Blind Frequency Offset Estimation*, Proceedings of IEEE International Workshop on Machine Learning for Signal Processing (Boston, MA), September 2015.
- [14] D. G. Tzikas, A. C. Likas, and N. P. Galatsanos, *The Variational Approximation for Bayesian Inference*, IEEE Signal Processing Magazine (2008), 131–146.
- [15] L. Wang and D. B. Dunson, *Fast Bayesian Inference in Dirichlet Process Mixture Models*, Journal of Computational and Graphical Statistics **20** (2011), no. 1, 196–216.